MNRAS 000, 1–12 (2016) Preprint 13 October 2016 Compiled using MNRAS LATEX style file v3.0

Resolving the blazar CGRaBS J0809+5341 in the presence of telescope systematics

Iniyan Natarajan[1]★, Zsolt Paragi[2], Jonathan Zwart[1,3], Simon Perkins[4], Oleg Smirnov[5,4], and Kurt van der Heyden[1]

[1]*Astrophysics, Cosmology and Gravity Centre (ACGC), Department of Astronomy, University of Cape Town, Private Bag X3, Rondebosch 7701, South Africa*
[2]*Joint Institute for VLBI ERIC, Postbus 2, 7990 AA Dwingeloo, the Netherlands*
[3]*Department of Physics and Astronomy, University of the Western Cape, Private Bag X17, Bellville 7535, South Africa*
[4]*SKA South Africa, 4th Floor, The Park, Park Road, Pinelands, 7405, South Africa*
[5]*Department of Physics and Electronics, Rhodes University, PO Box 94, Grahamstown 6140, South Africa*

arXiv:1610.03773v1 [astro-ph.IM] 12 Oct 2016

Last updated 2016 July 13; in original form 2016 July 13

ABSTRACT

We analyse Very Long Baseline Interferometry (VLBI) observations of the blazar CGRaBS J0809+5341 using Bayesian inference methods. The observation was carried out at 5 GHz using 8 telescopes that form part of the European VLBI Network. Imaging and deconvolution using traditional methods imply that the blazar is unresolved. To search for source structure beyond the diffraction limit, we perform Bayesian model selection between three source models (point, elliptical Gaussian, and circular Gaussian).Our modelling jointly accounts for antenna-dependent gains and system equivalent flux densities. We obtain posterior distributions for the various source and instrumental parameters along with the corresponding uncertainties and correlations between them. We find that there is very strong evidence ($>10^9$:1) in favour of elliptical Gaussian structure and using this model derive the apparent brightness temperature distribution of the blazar, accounting for uncertainties in the shape estimates. To test the integrity of our method, we also perform model selection on synthetic observations and use this to develop a Bayesian criterion for the minimum resolvable source size and consequently the maximum measurable brightness temperature for a given interferometer, dependent on the signal-to-noise ratio (SNR) of the data incorporating the aforementioned systematics. We find that calibration errors play an increasingly important role in determining the over-resolution limit for SNR$\gg$100. We show that it is possible to exploit the resolving power of future VLBI arrays down to about 5 per cent of the size of the naturally-weighted restoring beam, if the gain calibration is precise to 1 per cent or less.

Key words: methods: data analysis – methods: statistical – techniques: high angular resolution – techniques: interferometric – quasars: individual: CGRaBS J0809+5341

1 INTRODUCTION

The commonly used measure for the angular resolving power of an optical system is the *Rayleigh criterion* (Rayleigh 1879), which, in its simplest form, states that the minimum resolvable angle is proportional to the quantity λ/d, where λ is the wavelength of observation and d is the diameter of the aperture. For a given wavelength, the angular resolution can be improved by increasing the aperture size. The technique of *aperture synthesis* (Ryle 1955), in which multiple telescopes are operated together as a single instrument known as an *interferometer*, is employed in radio, optical, and infra-red astronomy to synthesise a large, partially-filled aperture in order to obtain high angular resolution. This is achieved by computing the *cross-correlation* of the radiation fields received at two different locations, also known as the *mutual coherence function*. In radio interferometry, the measured coherence function is a complex-valued *visibility*[1], and is related to the sky brightness distribution via a Fourier transform operation, under some simplifying assumptions such as the field-of-view being small enough (e.g., Wilson et al. 2009). By convention, this Fourier domain in which the visibilities are measured is known as the *visibility* domain or the *uv*-domain. Any two interferometer antennas (or *stations*) may be

★ Contact e-mail: iniyan@ast.uct.ac.za

[1] When the measurements are made using two mutually-orthogonal polarisation feeds that record the entire polarisation state of the radio wave, we obtain a *source coherency* or *visibility matrix* (e.g., Smirnov 2011). See also section 2.3.

© 2016 The Authors

said to form a *baseline*, and the length of the longest baseline for a given interferometer array configuration becomes the diameter d of the aperture. Thus, angular resolution of less than a milli-arcsecond (mas) can be achieved using Very Long Baseline Interferometry (VLBI) techniques in which data from radio telescopes separated by thousands of kilometres operating independently are synchronised post hoc using atomic clocks (e.g., Thompson et al. 2001; Middelberg & Bach 2008).

The theoretical limit for the resolution of an interferometer can be obtained using Fourier optics which characterises the resolving power of an optical system in terms of its *spatial bandwidth* (Goodman 1968). Diffraction effects limit the maximum spatial frequency, $f_{\rm max}$, transmitted by the optical system and the resolving power may be reformulated in terms of this frequency as $R = \pi/f_{max}$, where R is the *Nyquist distance* (Bertero & de Mol 1996). It is possible to obtain information about the spatial frequencies that lie outside f_{max} (i.e., beyond the *diffraction limit*) by incorporating *a priori* knowledge about the brightness distribution of the source. The resolving power then depends on the precision of the measuring instrument (e.g., Harris 1964; di Francia 1969). Introducing prior information to help solve a problem falls under the domain of inverse problem theory, which deals with the question of obtaining causes from results (Tarantola 2005). Most scientific questions are of this kind. More than one distribution of the relevant parameter values could have given rise to the same result[2] and hence it becomes necessary to impose additional constraints in order to obtain the most sensible explanation for the cause. This process is known as *regularisation* and is best understood in terms of probability theory. Under the Bayesian probability formalism, *a priori* information about a model is introduced in the form of prior distributions on its parameters.

Techniques for over-resolution or *super-resolution* have been widely used in radio interferometry, and especially in VLBI, to characterise partially resolved compact sources on milli-arcsecond scales. Lobanov (2005) derives the resolution limits for specific brightness-distribution templates for astronomical sources and Martí-Vidal et al. (2012) extend this to the general case of super-resolution with interferometers. A discussion of model-fitting in the *uv*-domain for specific source profiles using real and synthetic data using the UVMULTIFIT package is provided in Martí-Vidal et al. (2014). An alternative which achieves over-resolution by *sparse* modelling [3] in the sky or *image* domain in the presence of Gaussian noise generated with the same standard deviation on all baselines is explored in Honma et al. (2014). While these works explore the theoretical constraints of super-resolution, the application of these methods to real observations made with a specific interferometer configuration in the presence of instrumental uncertainties, will not provide us with a knowledge of how they correlate with the source parameters and limit the resolution. This issue will likely be relevant for snapshot VLBI observations using the Square Kilometre Array (SKA) (Paragi et al. 2015) as well.

Lochner et al. (2015) present an overview of the history and an implementation of Bayesian inference methods for simultaneous estimation of the sky and instrumental parameters through the analysis of visibilities. This work is organised as follows: In section 2, we explore the theoretical background for performing Bayesian analysis in the *uv*-domain and modelling visibilities using the measurement equation formalism. In section 3, we apply this approach to the European VLBI Network (EVN) observations of a flaring blazar CGRaBS J0809+5341 (henceforth, J0809+5341). The source is bright, and was observed for a relatively short time with a limited *uv*-coverage, to study possible structural changes associated with the flare (An et al. 2016). The situation is complicated by the fact that the station providing the long baselines in the array has no short spacings comparable to those provided by the rest of the stations, making its gain calibration difficult (Martí-Vidal et al. 2012). Our aim is to explore the resolving power of this interferometer and to better understand the uncertainties in the estimated source model parameters and how the station calibration errors affect these estimates. Section 4 presents our methodology in detail and section 5, our results. The same data were analysed using traditional methods by An et al. (2016). In section 6, we compare the minimum resolvable source sizes we obtain for this interferometer using the Bayesian approach on synthetic observations of compact sources with those predicted by Martí-Vidal et al. (2012), in the presence of station-dependent gains and baseline-dependent noise terms.

2 THEORETICAL BACKGROUND

2.1 Bayesian Inference

Statistical inference is the process of drawing conclusions about scientific propositions from data. A mathematically consistent way of doing this is by using Bayes' theorem to update our beliefs about propositions as more information becomes available (Jaynes 2003). In general, given two propositions A and B, and relevant prior information I, Bayes' theorem states that

$$\mathcal{P}(A|B,I) = \frac{\mathcal{P}(A|I)\,\mathcal{P}(B|A,I)}{\mathcal{P}(B|I)} \quad , \tag{1}$$

where $\mathcal{P}(X|I)$ denotes the conditional probability of proposition X given that information I is true and ',' denotes the conjunction *and*.

Bayesian inference can be performed at two levels (Mackay 2003). At the first level (*parameter estimation*), we assume that a model (or hypothesis) is true and fit its parameters to the data. If Θ denotes the set of parameters associated with the model H, then, given data D, Bayes' theorem may be rewritten as

$$\mathcal{P}(\Theta|D,H) = \frac{\mathcal{P}(\Theta|H)\,\mathcal{P}(D|\Theta,H)}{\mathcal{P}(D|H)} \quad . \tag{2}$$

$\mathcal{P}(\Theta|H) \equiv \Pi(\Theta)$ is called the *prior* probability distribution, which encodes our beliefs about the parameters prior to the analysis of the data[4]. $\mathcal{P}(\Theta|D,H)$ is the *posterior* probability distribution of the parameters which describes how the data D modify our initial beliefs. $\mathcal{P}(D|\Theta,H) \equiv \mathcal{L}(\Theta|D,H)$ is called the *likelihood*, which is a function of the parameters given the data. The mathematical form of the likelihood function reflects how the uncertainties in the measurements of the data are distributed (Trotta 2008).

The denominator in equation (2), $\mathcal{P}(D|H) \equiv \mathcal{Z}$, is called the *evidence* (or the *marginal likelihood*) and is obtained by integrating the numerator over $\Pi(\Theta)$:

$$\mathcal{Z} = \int \Pi(\Theta)\mathcal{L}(\Theta|D,H) \quad \mathrm{d}^N\Theta \quad , \tag{3}$$

[2] An *ill-posed* problem for which there is no unique solution.

[3] *Sparsity* is the ratio between the number of non-zero elements in a matrix and the total number of elements; a sparse matrix is one whose sparsity is high.

[4] The data may be captured before the prior probability distribution is drawn up. This is perfectly fine, since the prior occurs *logically* before the posterior in the chain of reasoning, not *temporally*.

Table 1. Criteria for model selection. B_{12} denotes the ratio of the evidences between hypotheses H_1 and H_2 (Kass & Raftery 1995).

$2\ln(B_{12})$	B_{12}	Evidence against H_2
0 to 2	1 to 3	Not worth more than a mention
2 to 6	3 to 20	Positive
6 to 10	20 to 150	Strong
> 10	> 150	Very strong

where N is the length of Θ. This process of *marginalisation* allows us to obtain the probabilities of a subset of the parameters without reference to the parameters with respect to which the integration is carried out. When the integration is performed over all of Θ, we get the evidence. This quantity may safely be ignored during parameter estimation (unless one intends to perform the second level of inference) since it does not affect the location or the shape of the posterior distribution in the parameter space and serves only as a normalisation constant.

The second level of inference (*model selection*) is determining the relative probabilities of alternative hypotheses given the data. Given hypothesis H, and a prior belief in the validity of H given by $\mathcal{P}(H|I)$, the model posterior probability may be computed using the evidence obtained from parameter estimation as

$$\mathcal{P}(H|D,I) \propto \mathcal{P}(D|H,I)\,\mathcal{P}(H|I) \quad . \tag{4}$$

The evidence is a quantitative measure of how much the data favour one model over another. It incorporates *Occam's razor*[5] automatically: unless a complicated model with more parameters (a high-dimensional parameter space) is significantly better at explaining the data (i.e., has a higher likelihood), its evidence will be smaller than that of a simpler model with fewer parameters that can explain the data equally well.

Given two models H_1 and H_2, we may define a model selection ratio between the posteriors of the two models as

$$\frac{\mathcal{P}(H_1|D,I)}{\mathcal{P}(H_2|D,I)} = \frac{\mathcal{Z}_1}{\mathcal{Z}_2}\,\frac{\mathcal{P}(H_1|I)}{\mathcal{P}(H_2|I)} = B_{12}\,\frac{\mathcal{P}(H_1|I)}{\mathcal{P}(H_2|I)} \quad , \tag{5}$$

where $\mathcal{P}(H_1|I)/\mathcal{P}(H_2|I)$ is the ratio of the prior probability distributions of the two models which may often be set to unity, indicating that there is no prior preference for one model over the other. The ratio of the evidences, B_{12}, is known as the *Bayes factor* (Jeffreys 1961). The higher this factor is, the more is H_1 preferred over H_2. It is useful to consider the natural logarithm of the Bayes factor:

$$\ln(B_{12}) = \ln(\mathcal{Z}_1) - \ln(\mathcal{Z}_2) \quad . \tag{6}$$

We use twice this value as a measure of how strongly a model is preferred over another (Kass & Raftery 1995) (Table 1).

Frequently, the distributions involved are analytically intractable. Numerical techniques such as MCMC (Markov Chain Monte Carlo) (Metropolis et al. 1953; Hastings 1970) are used to sample a multi-dimensional parameter space, resulting in the estimation of the joint posterior distribution of the parameters, but they are inefficient when required to compute the Bayesian evidence, since it is a multi-dimensional integral. We therefore use the MultiNest[6] algorithm (Feroz & Hobson 2008; Feroz et al. 2009, 2013), based on the nested sampling method originally proposed by Skilling (2004), a Monte Carlo method efficient at sampling posteriors with multiple modes and/or curving degeneracies in low-dimensional parameter spaces ($\lesssim$ 30). MultiNest also computes the evidence at a fraction of the cost one would incur with the standard MCMC techniques (Buchner et al. 2014).

[5] Occam's razor states that, among competing hypotheses, the one that makes the fewest assumptions should be favoured.

[6] https://ccpforge.cse.rl.ac.uk/gf/project/multinest

2.2 Bayesian Analysis of Visibilities

Estimating the true brightness distribution of the sky from radio data falls under the class of *inverse problems*, which aim to determine the underlying phenomena (the causes) from the observed data (the results) (Parker 1977). The technique of *model-fitting* is generally suited to handling such problems. It is most useful when the source brightness distribution and the instrumental effects can be accurately represented in parametric form, giving reliable estimates that cannot be obtained from a deconvolved image alone. We devise a model with adjustable parameters, which we believe is capable of describing the data, and choose a *figure-of-merit* or *merit function* to measure how well the data and the model agree with each other (Press et al. 2007, Chapter 15). In the Bayesian approach, the merit function is just the posterior probability distribution and the data, in the present case, are the observed visibilities.

Statistical analysis of visibilities complements the traditional imaging and deconvolution techniques and, if applied judiciously, can improve on them. Sparse *uv*-coverage, the spreading of localised *uv*-domain errors throughout the image (thereby making them correlated between pixels) by the Fourier transform process, and the difficulty in estimating the measurement uncertainties (or the *noise*) from an image deconvolved using non-linear deconvolution techniques such as CLEAN (Högbom 1974a) or MEM (Ables 1974), are but some of the factors that render Fourier-transform imaging difficult and often inadequate (Taylor et al. 1999, Chapter 16).

Our primary assumptions are that (i) the visibility measurements are independent and (ii) the uncertainties in the measurements follow a Gaussian distribution. If the data are independent, then the likelihood $\mathcal{L}$ from equation (2) can be expressed as (Sivia & Skilling 2006, Chapter 3)

$$\mathcal{P}(D|\Theta,H) = \prod_{k=1}^{N_{\rm vis}} \mathcal{P}(D_k|\Theta,H) \quad , \tag{7}$$

where D_k stands for the kth datum and $N_{\rm vis}$ denotes the number of visibilities. Under the assumption that the visibility noise is Gaussian, the probability of obtaining each individual datum is given by

$$\mathcal{P}(D_k|\Theta,H) = \frac{1}{\sigma_k\sqrt{2\pi}} \exp\left(-\frac{(F_k - D_k)^2}{2\sigma_k^2}\right) \quad , \tag{8}$$

where $\{\sigma_k\}$ denotes the expected uncertainties and $F_k \equiv f(\Theta,k)$ denotes a functional form of the model for a specific set of parameter values Θ for the $k^{\rm th}$ datum.

Combining equations (7) and (8), we obtain

$$\mathcal{P}(D|\Theta,H) \propto \exp\left(-\frac{\chi^2}{2}\right) \quad , \tag{9}$$

where the sum of the squares of the *normalised residuals* is represented by

$$\chi^2 = \sum_{k=1}^{N}\left(\frac{F_k - D_k}{\sigma_k}\right)^2 \quad . \tag{10}$$

If we also assume a flat (constant) prior for the parameters to indicate that we are largely ignorant of their expected values, then,

taking the natural logarithm of equation (2) and omitting the normalisation constant, we arrive at

$$\ln[\mathcal{P}(\Theta|D,H)] = \mathrm{C}(\sigma_k) - \frac{\chi^2}{2} \quad . \tag{11}$$

where $\mathrm{C}(\sigma_k)$ depends only on σ_k. Thus, with some simplifying assumptions, Bayesian parameter estimation reduces to the more familiar methods of *maximum likelihood* and *least-squares estimation*. The crucial difference is that Bayesian inference is independent of the Gaussianity of the underlying process. The exact form our likelihood function takes is very close to this equation (section 4.3).

2.3 The Measurement Equation

The simultaneous estimation of source and instrumental parameters is facilitated by the *Radio Interferometer Measurement Equation* (RIME), originally formulated by Hamaker et al. (1996) and reformulated for *direction-dependent effects* (DDEs) by Smirnov (2011). The RIME unifies the concept of the Stokes parameter representation of electromagnetic waves (Born & Wolf 1999) and the technique of radio interferometry under a single mathematical framework.

For a sky composed of discrete point sources, the RIME may be written as

$$\mathbf{V}_{pq} = \sum_s \boldsymbol{J}_{sp}\, \mathbf{B}_s\, \boldsymbol{J}^H_{sq} \quad , \tag{12}$$

where $\mathbf{V}_{pq}$ denotes the 2×2 matrix of visibilities measured by the baseline pq with contribution from each discrete source s, $\boldsymbol{J}_{sp}$ is a *Jones matrix* that incorporates all propagation path effects on the way from the source s to the antenna p, the superscript H denotes the Hermitian conjugate, and $\mathbf{B}_s$ is the brightness matrix of source s, which relates the correlated voltages from the two polarisation feeds to the four Stokes parameters I, Q, U, and V[7]:

$$\mathbf{B}_s = \begin{pmatrix} I+V & Q+iU \\ Q-iU & I-V \end{pmatrix} \quad . \tag{13}$$

In the more general case, if we consider the sky to be a continuous brightness distribution, the RIME may be written as

$$\mathbf{V}_{pq} = \int_{4\pi} \boldsymbol{J}_p(\widehat{\boldsymbol{\sigma}})\, \mathbf{B}(\widehat{\boldsymbol{\sigma}})\, \boldsymbol{J}^H_q(\widehat{\boldsymbol{\sigma}}) \quad \mathrm{d}\Omega \quad , \tag{14}$$

where $\widehat{\boldsymbol{\sigma}}$ represents the unit direction vector and $d\Omega$ denotes spherical integration over the entire sky. The known DDEs and DIEs (*direction-independent effects*) may each be assigned its own Jones matrix, while the unknown effects may be subsumed into a generic Jones matrix term. This approach is called the *phenomenological* RIME and has been implemented in the MeqTrees[8] software suite for interferometric simulation and calibration (Noordam & Smirnov 2010) that we use for our simulations (section 6). To interface with MultiNest, which samples the posteriors and computes the evidence, we use PyMultiNest[9], a python wrapper to MultiNest (Buchner et al. 2014).

Unlike the toy model selection problem presented in Lochner et al. (2015), the models we use here incorporate instrumental systematics and hence contain more parameters (section 4.1). To evaluate the likelihood and consequently, estimate the evidence, the RIME must be computed at each iteration of the Bayesian inference process. To accurately model the visibilities given a sparse sky model, the RIME is implemented using the computationally expensive direct Fourier transform. We use *Montblanc*[10] (Perkins et al. 2015), a GPU-accelerated implementation of the RIME for this. Montblanc is built on PyCUDA, a Python interface to NVIDIA's CUDA architecture (Klöckner et al. 2012), to parallelise the RIME evaluation. Montblanc computes the RIME from sampled model parameters during each iteration to generate the χ^2 value (equation 10), which may be used for calculating the likelihood. Currently, it supports the simulation of three source morphologies (point, Gaussian, and Sérsic) of which we use the point and Gaussian source models, and the inclusion of DIEs and DDEs in the RIME.

Table 2. EVN stations used in the observation along with the corresponding dish diameters and the nominal (SEFD) values.

Station	Code	Diameter (m)	Nominal SEFD (Jy)
Effelsberg	EF	100	20
Jodrell Bank	JB	25	320
Noto	NT	32	260
Onsala	ON	25	600
Torun	TR	32	220
OAN-Yebes	YS	40	160
Westerbork	WB	25	120
Sheshan	SH	25	720

3 DATA DESCRIPTION AND REDUCTION

Blazars are galaxies that exhibit strong radio emission from their cores (called *active galactic nuclei* (AGN)) in the form of jets of relativistically beamed particles along the line of sight to the observer (e.g., Antonucci 1993; Urry & Padovani 1995). They display high radio luminosities (10^{46} erg s^{-1}) generated by non-thermal radiation mechanisms, mainly synchrotron and inverse Compton radiation (Rybicki & Lightman 1985), with flat spectra. Owing to their compactness, blazars are an important class of objects whose studies have largely benefited from VLBI observations.

J0809+5341 is located at J$08^{\mathrm{h}}09^{\mathrm{m}}41.733^{\mathrm{s}}$, $+53^{\mathrm{d}}41^{\mathrm{m}}25.092^{\mathrm{s}}$ at redshift $z = 2.144$ (Pâris et al. 2014). It was observed using 8 EVN stations (Table 2), for 130 minutes using 8×16 MHz spectral bands of 32 channels each, between 4926 and 5054 MHz, with 2 s integration time, in both R and L polarisations. The *uv*-coverage is shown in Figure 1. The long baselines correspond to the Sheshan station in Shanghai (SH) and extend over 9000 km. Without SH, the maximum baseline length is about 2200 km. More details on this project can be found in An et al. (2016).

Invalid data were excluded (or *flagged*) before the visibilities were *self-calibrated* (Readhead & Wilkinson 1978)[11]. Self-calibration is the process of introducing some plausible assumptions about the source structure to correct the observed complex visibilities (Taylor et al. 1999). For this procedure, the channels in each band were averaged together so that the data contained 8 bands

[7] Here, we assume that the two feeds are sensitive to circular polarisation; for linear polarisation, the brightness matrix undergoes a further linear transformation (Smirnov 2011, section 6.3).
[8] http://meqtrees.net
[9] https://johannesbuchner.github.io/PyMultiNest
[10] https://github.com/ska-sa/montblanc
[11] www.evlbi.org/user_guide/evn_datareduc.html

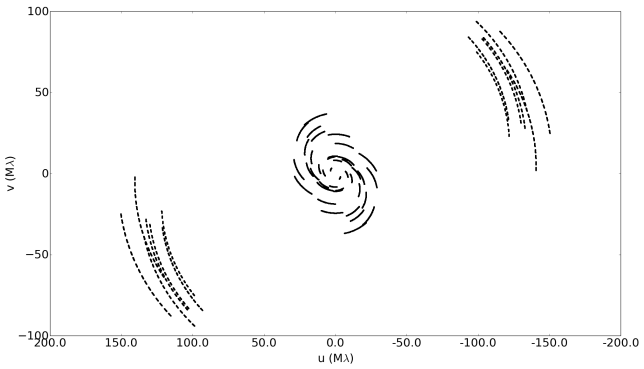

Figure 1. The *uv*-coverage of the synthesis.

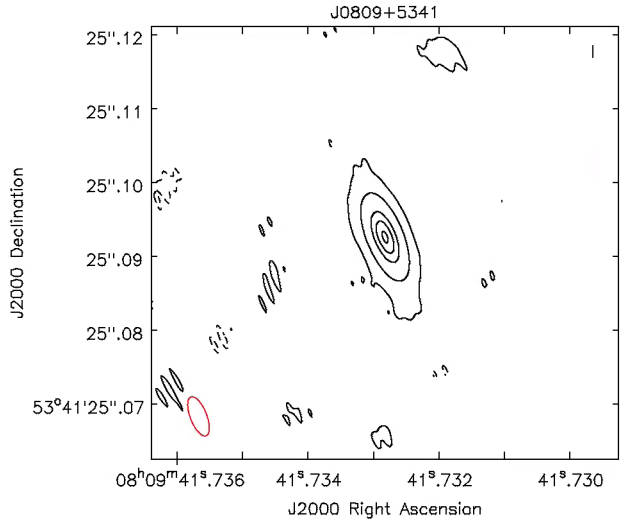

Figure 2. Stokes *I* contour image of J0809+5341. The contours presented are $-3, 3, 100, 1000, 2000$, and 3000 times the rms noise in the image ($\sim$ 0.05 mJy), with the negative contours shown as dashed lines. The red ellipse at the bottom left corner is the PSF used for restoration by CLEAN: 5.7×2.2 mas, oriented at an angle of 21.7°.

of one channel each. This frequency-averaging is made possible by the fact that the source is located at the pointing centre and is therefore not affected by frequency smearing (Smirnov 2011). We extracted the spectral band centred at 4982.24 MHz for our analysis. The original UVFITS file was converted into the Measurement Set (MS) format and the missing[12] baselines were introduced – with the corresponding data flagged – for compatibility with Montblanc.

A *naturally*-weighted (Briggs 1995), deconvolved image of the self-calibrated data shows a compact source (Figure 2). The deconvolution was performed using the CSCLEAN algorithm (Schwab 1984) in *lwimager*[13] with 1000 iterations. The point spread function (PSF) of the interferometer was calculated using the Hogbom algorithm recommended for data with poor *uv*-coverage (Högbom 1974b). The dynamic range of the image is 3200:1. Image-plane source extraction using PyBDSM[14] estimates that the source is unresolved with a flux density of 161.1 ± 0.01 mJy. In comparison, An et al. (2016) performed model-fitting on the data using *Difmap*[15], a software package for self-calibrating visibilities and iteratively building up a model of the sky (Shepherd 1997), and concluded that the best-fit model was a circular Gaussian of size 0.31 ± 0.06 mas.

[12] Due to the source not being visible to those baselines.
[13] https://github.com/casacore/casarest
[14] http://www.astron.nl/citt/pybdsm

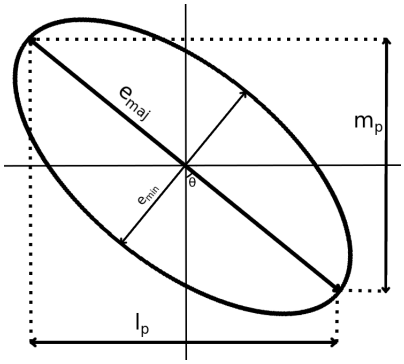

Figure 3. The relation between the shape parameters, l_p, m_p, and r, and the major axis ($e_{\rm maj}$), minor axis ($e_{\rm min}$), and position angle (θ) of an elliptical Gaussian profile.

4 BAYESIAN ANALYSIS

4.1 Description of Models

We compare three morphological models to describe unresolved or partially-resolved sources, each differing in the assumed brightness distribution of the source:

PT: The point source model consists of one flux density parameter, S_ν, and two parameters that describe the position of the source, (l, m), the direction cosines measured with respect to the (u, v) coordinates.

GAU: The elliptical Gaussian model consists of S_ν and (l, m) along with three more parameters, l_p, m_p, and r, describing the shape of the source. These three parameters are related to the major axis $e_{\rm maj}$, minor axis $e_{\rm min}$, and position angle or orientation θ of the ellipse as follows:

$$\begin{aligned} l_p &= e_{\rm maj} \sin\theta \\ m_p &= e_{\rm maj} \cos\theta \\ r &= e_{\rm min}/e_{\rm maj} \quad . \end{aligned} \tag{15}$$

Thus l_p and m_p denote the projections of $e_{\rm maj}$, onto the l and m axes, and r is the ratio of the minor axis to the major axis (see Figure 3).

CIRC: The circular Gaussian model consists of the same six parameters as GAU, with the exception that the ratio r is given a delta distribution at unity, forcing the shape to be circular.

A summary of the parametrisation of each model is given in Table 3. Alongside the parameters that describe the source, all three models incorporate the same instrumental parameters. In principle, any instrumental effect could be modelled using the RIME (see section 2.3), while in practice, we are often limited by the software used to sample the parameter space. For instance, MULTINEST performs best with low-dimensional ($\lesssim 30$) parameter spaces. For the present experiment, the most important parameters that affect the

[15] ftp://ftp.astro.caltech.edu/pub/difmap/difmap.html

Table 3. Models evaluated in this study. Besides the source parameters, there are 7 free parameters describing the station gain amplitudes and 8 parameters describing the individual station SEFDs.

Model	Number of parameters	Parametrisation
PT	18	Flux Density (S_ν) Position (l, m) Station gain amplitudes ($\|g_p\|$) SEFD$_p$
GAU	21	Flux Density (S_ν) Position (l, m) Shape (l_p, m_p, r) Station gain amplitudes ($\|g_p\|$) SEFD$_p$
CIRC	20	Flux Density (S_ν) Position (l, m) Shape (l_p, m_p) Station gain amplitudes ($\|g_p\|$) SEFD$_p$

source shape resolution are the amplitudes of the complex station gains and the *System Equivalent Flux Densities* (SEFD) of each antenna. The complex station gains vary with respect to time during the course of the observation and are solved for during the preliminary self-calibration. The resulting gain amplitudes and phases are sufficiently smooth to be modelled using one complex gain term per station. Moreover, since all the source models considered place the source at the pointing centre, the measured phases on all the baselines will be zero. Hence, we hold the gain phases constant at zero. The SEFD is the flux density equivalent of the system temperature T_{sys} of a receiver system (Thompson et al. 2001). It varies widely between stations in any VLBI observation and is here used to derive the per-visibility noise term σ_{pq} for each baseline pq, which are then used to weight the model visibilities using the inverse of the corresponding variance. This brings down the number of parameters that describe the noise from the *number of baselines* to the *number of stations*. The per-visibility uncertainty for one polarisation in terms of the geometric mean of the station SEFDs is given by the *radiometer equation* (Thompson et al. 2001):

$$\sigma_{pq} = \frac{\text{SEFD}_{pq}}{\sqrt{\delta\nu\,\tau_{pq}}} \quad , \qquad \text{where SEFD}_{pq} = \sqrt{\text{SEFD}_p \,\text{x}\, \text{SEFD}_q} \quad , \tag{16}$$

SEFD$_p$ is the SEFD of station p, $\delta\nu$ is the channel bandwidth, and τ_{pq} is the integration time for baseline pq.

4.2 Building the RIME

We are now ready to construct the RIME for modelling the visibilities. Assuming a flat spectral index and using the flux density parameter S_ν in the brightness matrix (equation 13), we obtain

$$\mathbf{B} = \begin{pmatrix} S_\nu & 0 \\ 0 & S_\nu \end{pmatrix} \quad . \tag{17}$$

The first linear transformation this signal undergoes is represented by the phase delay matrix K, associated with the difference in the geometric path lengths from the source to antennas p and q. Given the phase difference (κ_p) between the waves received by antenna p located at $\mathbf{u}_p = (u_p, v_p, w_p)$ relative to $\mathbf{u} = 0$, the scalar *K-Jones* term for antenna p can be written as

$$K_p = \mathrm{e}^{-i\kappa_p} \equiv \mathrm{e}^{-i\kappa_p} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad . \tag{18}$$

The K-Jones term must be accounted for even under ideal conditions in which nothing else affects the signal from the source to the interferometer. Knowing this, we may define the *source coherency* matrix as (Smirnov 2011)

$$\mathbf{X}_{pq} = K_p\, \mathbf{B}\, K_q^H = \mathbf{B}\mathrm{e}^{-i\kappa_{pq}} \quad . \tag{19}$$

If we assume equal gains for the two polarisation feeds, then $g_p = g_{x_p} = g_{y_p}$, and the diagonal *G-Jones* matrix describing station gains reduces to a scalar matrix,

$$G_p = \begin{pmatrix} g_p & 0 \\ 0 & g_p \end{pmatrix} = g_p \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad . \tag{20}$$

The PSF of the observation is 5.7×2.2 mas when the visibilities are weighted naturally, while the primary beams of the stations are at least 8 arcmin wide at 5 GHz. Hence, for a source located at the pointing centre, the *E-Jones* matrix for primary beam effects may be set to unity, further simplifying the RIME.

Now, remembering the additive Gaussian noise term with zero mean and a variance of σ^2_{pq} per visibility, the RIME for the point source becomes

$$\mathbf{V}_{pq} = G_p\, \mathbf{X}_{pq}\, G_q^H + \mathcal{N}(0, \sigma^2_{pq}) \quad . \tag{21}$$

For the extended source models, the brightness distribution is integrated over the extent of the source. Expressed in terms of the direction cosines l and m, the RIME for the extended source may be written as (Thompson et al. 2001, section 3.1)

$$\mathbf{V}_{pq} = \iint_{lm} G_p\, \mathbf{X}_{pq}(l, m)\, G_q^H \, \mathrm{d}\Omega \; + \; \mathcal{N}(0, \sigma^2_{pq}), \qquad \text{where } \mathrm{d}\Omega = \frac{\mathrm{d}l\, \mathrm{d}m}{\sqrt{1 - l^2 - m^2}} \quad . \tag{22}$$

4.3 Likelihood function

Now that we have described the models quantitatively, we can set up the likelihood function for the problem at hand. Following the discussion in section 2.2, given the observed (V_D) and the modelled (V_M) visibilities, and the uncertainties σ_k that vary with baseline, the likelihood function for parameter estimation for model H may be written as

$$\mathcal{L}(\Theta|V_D, H) = \frac{1}{\prod_{k=1}^{2N_{\text{vis}}} \sqrt{2\pi\sigma_k^2}} \exp\left(-\frac{\chi^2}{2}\right) \quad , \qquad \text{where } \chi^2 = \sum_{k=1}^{2N_{\text{vis}}} \left(\frac{V_{M_k} - V_{D_k}}{\sigma_k}\right)^2 \quad , \tag{23}$$

and N_{vis} is the total number of complex visibilities. The summation is carried out over $2N_{\text{vis}}$ since we consider the real and imaginary parts separately. Taking the natural logarithm of $\mathcal{L}$, we obtain

$$\begin{aligned} \ln(\mathcal{L}) &= \sum_{k=1}^{2N_{\text{vis}}} \ln\left[(2\pi\sigma_k^2)^{-1/2}\right] - \frac{\chi^2}{2} \\ &= -\frac{1}{2} \sum_{k=1}^{2N_{\text{vis}}} \ln\left[2\pi\sigma_k^2\right] - \frac{\chi^2}{2} \quad . \end{aligned} \tag{24}$$

Table 4. Prior distributions for the different parameters used. All the listed parameters were set uniform priors with the range indicated by the values in the square brackets. The parameters with delta priors are not included.

Parameter (units)	Prior distribution
S_ν / Jy	[0.1, 0.2]
l&m / mas	[-4, 4]
l_p / mas	[0, +4]
m_p / mas	[-4, +4]
r	[0, 1]
$\|g_p\|$, where $p \neq$ EF	[0.8, 1.2]
SEFD / Jy	[5, 800]

The noise is modelled such that the variance for the real and imaginary parts of a complex visibility is the same. Hence, counting each σ_k twice, we arrive at the final form of the log-likelihood function:

$$\ln(\mathcal{L}) = -\sum_{k=1}^{N_{\rm vis}} \ln\left[2\pi\sigma_k^2\right] - \frac{\chi^2}{2} \quad . \tag{25}$$

4.4 Prior distributions

We set uniform (flat) priors $\Pi(\Theta)$ on most parameters (Table 4). The prior range for S_ν is chosen based on what we know about the flux density of J0809+5341 from preliminary imaging and source extraction (section 3). The prior distribution for l_p (the sine projection of $e_{\rm maj}$) is restricted to non-negative values so that the position angle estimate is constrained to a range of 180°. We allow the gain amplitudes to vary between ±20 per cent of unity and keep the corresponding phases fixed at zero[16]. To break the degeneracy between S_ν and $|g_p|$ estimates, the gain amplitude of the EF station is set a delta prior at unity; we choose EF because it provides some of the shortest baselines (except for the baseline with SH) in the synthesis and is not sensitive to the structure of the compact source, and the preliminary self-calibration tells us that EF has the most stable gain.

For the model selection step, we assign equal priors to all the models considered, so that the logarithm of the Bayes factor B_{12} may directly be used for model comparison.

5 RESULTS

The MultiNest sampler takes the prior distributions and the likelihood as its inputs, computes the natural logarithm of the Bayesian evidence (equation (3)) for each model, and produces the joint posterior as a by-product. We evaluate each model independently and use the corresponding logarithmic evidences to compute the Bayes factor between models using equation (6).

The analysis reveals that the factors $2\ln(B_{ij})$, for GAU against PT and CIRC are 21.0 ± 0.8 and 26.0 ± 0.8 respectively. Applying the criteria we set aside in Table 1, we conclude that there is a *very strong* preference for the elliptical Gaussian (GAU) over the other two models[17].

Figure 4 shows the correlations between various parameters

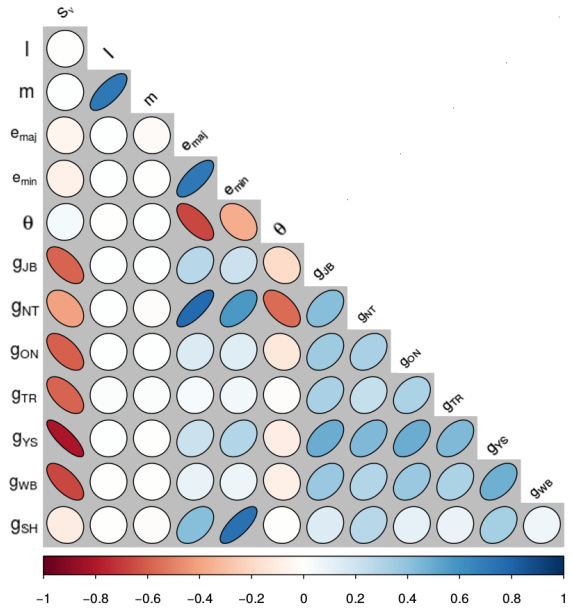

Figure 4. Correlations between the estimated parameters for model GAU. The quantities $e_{\rm maj}$, $e_{\rm min}$, and θ are derived from the shape parameters. The individual station SEFDs and parameters with delta prior distributions are excluded from the figure. The coloured ellipses correspond to the correlation coefficients shown at the bottom; the higher the ellipticity, the stronger the correlation.

of GAU. We use the quantities $e_{\rm maj}$, $e_{\rm min}$, and θ, obtained from equations (15), for visualisation. The individual station SEFDs are mostly uncorrelated with the other parameters and with each other, and are not shown in the figure. The station gain amplitudes are correlated negatively with the estimated flux density as one would expect: the higher the instrument gain, the lower the true flux density of the source. The gain amplitude, $|g_{\rm SH}|$, of SH, which provides the longest baselines, is correlated positively with the shape parameter estimates.

Figure 5 shows the 1-D and 2-D marginalised posteriors of the source and the station gain parameters, along with the source brightness temperature, T_b (section 5.3). The 2-D marginalised posterior distributions between $|g_{\rm SH}|$ and the three shape parameters show the precise nature of the relationships between them. Presenting these relationships in full is the most complete statement we can make about these parameters. The SH gain amplitude is not constrained to the same precision as those of the other stations by the preliminary self-calibration process. SH does not form short baselines with any other station, thereby making it difficult for amplitude self-calibration to correct for the biased gains (Martí-Vidal et al. 2012). Hence, measurements involving SH has contributions both from the resolved source structure and the uncertainties in the calibration of $|g_{\rm SH}|$. Figure 6 highlights the regions in the marginalised posteriors of $e_{\rm maj}$ and $e_{\rm min}$ that correspond to a 1 per cent variation in that of $|g_{\rm SH}|$ (which varies by about 12 per cent) about its mean. Providing short baselines for outlying stations such as SH by introducing more VLBI antennas close by would reduce the sparsity of the array distribution and minimise the uncertainty in the estimation of the station gains. In Shanghai, there are now two radio telescopes available for VLBI: the old Sheshan 25m (SH) and

[16] This is achieved by setting delta distributions centred at the known values of the parameters. The data do not impact the posteriors of such parameters since their posteriors are also delta distributions at the same locations in the parameter space.

[17] The quoted relative ln-evidences are obtained from the *importance nested sampling* results owing to their better accuracy, while the uncertainties are obtained from the *vanilla* nested sampling results (Feroz et al. 2013).

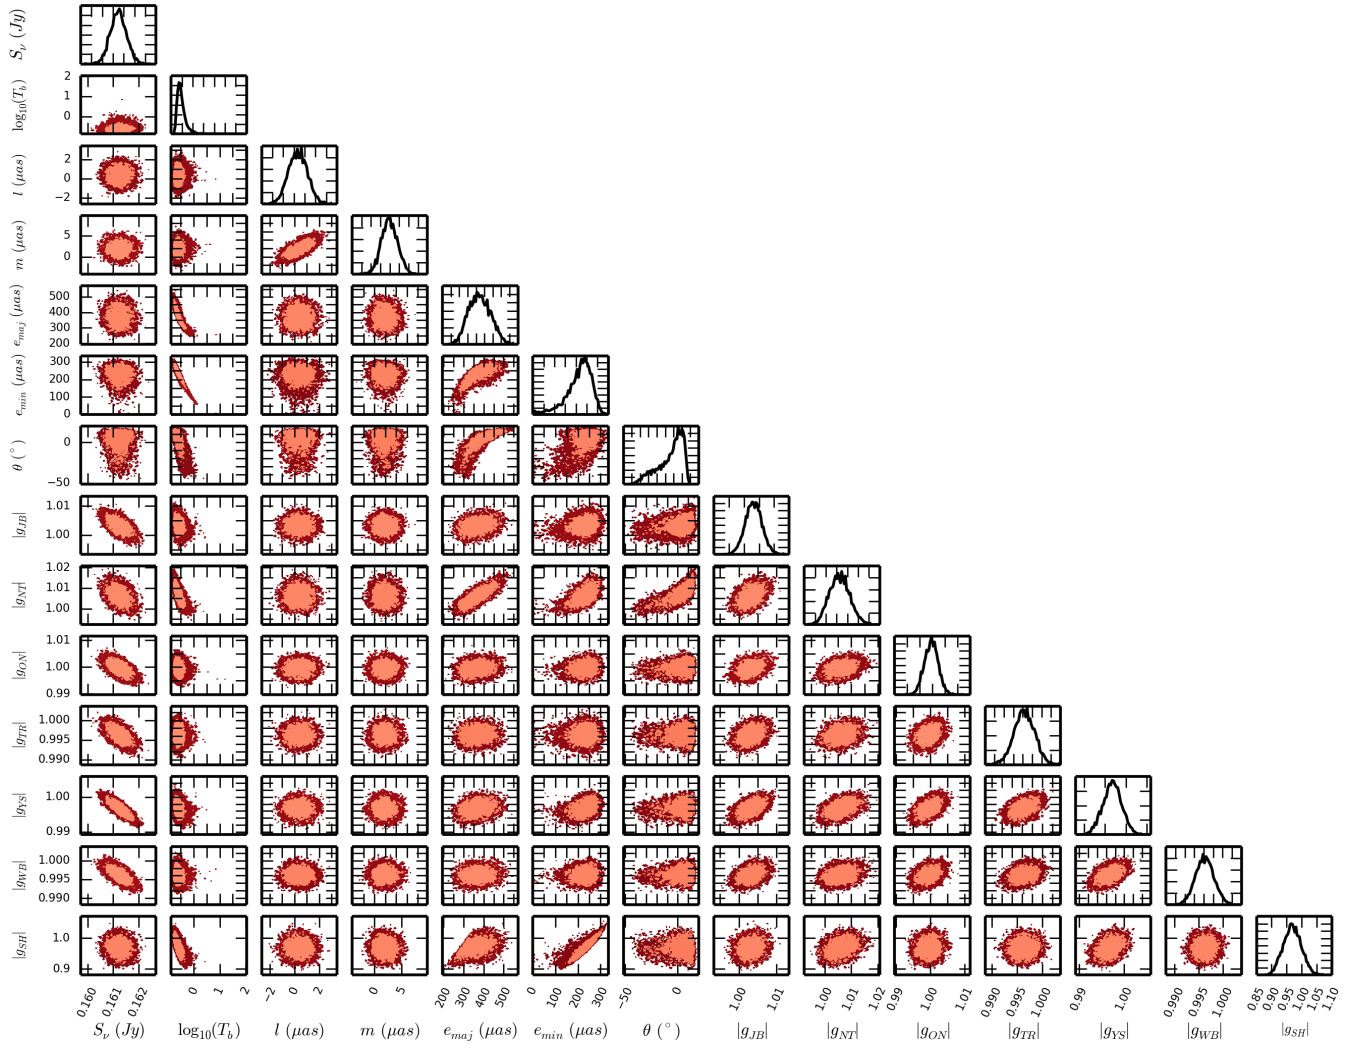

Figure 5. The 1-D posteriors and 2-D correlations of the source and the station gain parameters for model GAU. The principal diagonal gives the 1-D marginalised posterior distributions of the estimated flux density (S_ν), the source brightness temperature (T_b), the derived shape quantities ($e_{\rm maj}$, $e_{\rm min}$, θ), and the station gain amplitudes $|g_p|$, while the lower triangular matrix gives the 2-D joint posteriors between their various combinations. The 68 and 95 per cent credible regions are indicated by the light-red and dark-red shaded regions respectively. Parameters with delta priors and the station SEFDs (which are uncorrelated with the other parameters) are excluded from the figure.

the new Tianma 65m telescope which provides the necessary short baselines for SH (Kawaguchi et al. 2015). When both are used in VLBI observations, the gain calibration of the longest EVN baselines will improve and result in reduced uncertainty in the derived shape parameter and brightness temperature distributions.

There is also a strong correlation between the position estimates (although with uncertainties of the order of μas), which may be due to the position not being constrained tightly enough along the extent of the major axis of the PSF.

5.1 European stations only

The NT gain ampltude is positively correlated with the shape parameters (Figures 4 & 5). In the absence of SH, NT provides the longest baselines for the observation. This correlation motivated us to test whether the source can successfully be resolved using only the baselines pertaining to the European stations. The relative evidence between PT and GAU comes out to 2.86 ± 0.75, while that between PT and CIRC is 1.61 ± 0.76. There is positive evidence for PT and a mild preference for CIRC over GAU and so, without the SH measurements, we are forced to conclude that the source is barely resolved.

5.2 Comparison with DIFMAP results

We also performed model-fitting to the data using Difmap, conventionally used in VLBI, so that we could compare our results with those reported by An et al. (2016). Difmap returns the best-fit parameter estimates and the reduced chi-squared, $\chi^2_{\rm red}$, for each model as a measure of its goodness-of-fit[18]. The $\chi^2_{\rm red}$ values for these models, with and without the SH baselines, are shown in Table 5. These

[18] $\chi^2_{\rm red} = \chi^2$ / DoF, where DoF stands for *degrees of freedom*, obtained by subtracting the number of model parameters from the number of measurements.

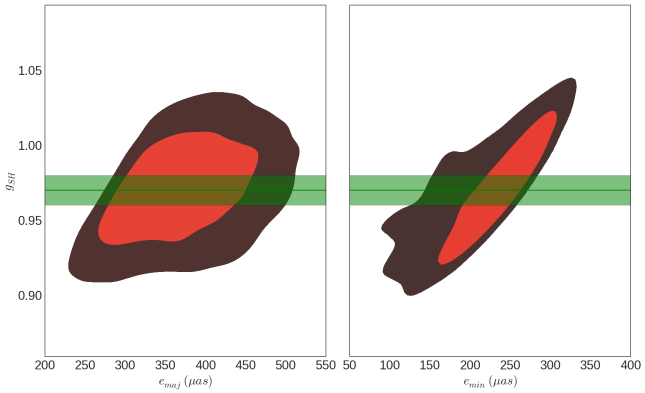

Figure 6. The correlations between SH gain amplitude and the shape parameters. The horizontal bars on either side of the regions shaded green correspond to ±1 per cent about the mean of the marginalised posterior of $g_{\rm SH}$. The light and dark red shaded regions indicate the 68 and 95 per cent credible regions respectively of the posteriors of the shape parameters $e_{\rm maj}$ and $e_{\rm min}$. It is evident that a better gain calibration would constrain both $e_{\rm min}$ and $e_{\rm maj}$, due to the strong correlation between $e_{\rm maj}$ and $e_{\rm min}$ (not highlighted here), better.

Table 5. Reduced chi-squared values for the three models with and without the SH measurements using Difmap. The asterisk indicates that these models do not take instrumental effects into account.

Model	$\chi^2_{\rm red}$ with SH (DoF = 164173)	$\chi^2_{\rm red}$ without SH (DoF = 121501)
PT*	2.2442	1.7266
GAU*	1.8746	1.7253
CIRC*	1.8785	1.7256

models include only the source parameters and do not account for the instrumental effects considered in the original three models (Table 3). As a result, they are much simpler than the models we test using the RIME.

With the SH measurements included, there is a slight preference for the resolved source models. Without the SH measurements, as with the Bayesian analysis, PT is at least as good a fit as any other model if not better, because of its simplicity. Though $\chi^2_{\rm red}$ has traditionally been used for model selection, it is a poor substitute for Bayesian evidence since it assumes that the underlying processes are Gaussian. The Bayesian approach does not depend on this assumption. Moreover, there is no measurement of the uncertainty in the value of $\chi^2_{\rm red}$, which becomes important when the $\chi^2_{\rm red}$ values of two models do not differ much. In contrast, the Bayesian evidence has an uncertainty associated with it which we may use to determine the significance of the model selection ratio.

Figure 7 shows the relation between the Difmap best-estimates and the posterior distributions of the source parameters obtained using the Bayesian analysis. The Difmap estimates are point estimates with no associated uncertainties. They fall within the 68 per cent credible region of the Bayesian posteriors when the posteriors are Gaussian (S_ν and $e_{\rm maj}$). For $e_{\rm min}$ and θ, the posteriors are non-Gaussian and must be presented and accounted for in full while drawing inferences, such as computing the brightness temperature distribution of the blazar.

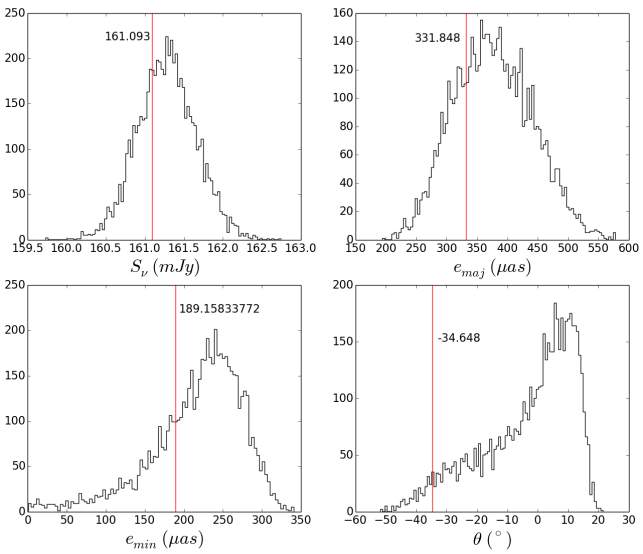

Figure 7. Comparison of the posteriors of the source parameters with the corresponding Difmap estimates. The vertical red lines correspond to the best-fit Difmap values printed alongside and the histograms correspond to the 1-D marginalised posteriors of the source parameters.

5.3 What does this mean for the blazar?

In the radio regime, the brightness temperature T_b of a source is given by the Rayleigh-Jeans approximation to Planck's law as (Wilson et al. 2009)

$$T_b = S_\nu \frac{c^2}{2k\nu^2} \frac{1}{\Delta\Omega} \quad , \tag{26}$$

where ν is the frequency of observation, k is the Boltzmann constant, and $\Delta\Omega$ is the solid area subtended by the source. For a Gaussian source at high redshifts, such as J0809+5341, this relation becomes (Kellerman & Owen 1988)

$$T_b = 1.22 \frac{S_\nu}{\nu^2\, e_{\rm maj} e_{\rm min}} (1+z)\ 10^{12}\ {\rm K} \quad , \tag{27}$$

With $\nu = 4.98224$ GHz and $z = 2.144$, we use the full posterior distributions of S_ν, $e_{\rm maj}$, and $e_{\rm min}$ to infer the distribution of T_b (Figure 8). This ensures that our measurements account for the uncertainties in the source shape. The mode of the distribution occurs at 0.25×10^{12} K. The 68 per cent credible region places T_b between 0.17×10^{12} K and 0.36×10^{12} K. The few very high values of T_b correspond to the very low values of $e_{\rm min}$.

Rees (1966) originally proposed that bright compact objects at cosmological distances are composed of parts moving at relativistic speeds. This bulk relativistic motion is the most probable model that explains the high brightness temperatures and apparent superluminal motion of jets observed in radio galaxies (e.g., Marscher & Scott 1980). An inverse Compton limit of $\simeq 10^{11}$–10^{12} K for the brightness temperature was derived by Kellerman & Pauliny-Toth (1969), which is the range in which our T_b measurements for the blazar lie. This enhanced brightness temperature is a consequence of *Doppler boosting* in which the relativistic beaming of a jet moving at an acute angle to the line of sight to the observer increases the observed flux density without increasing the size of the jet (Kellerman & Owen 1988). The measured brightness temperature and the internal brightness temperature, $T_{b,int}$, relate as $T_b = \delta\, T_{b,int}$, where

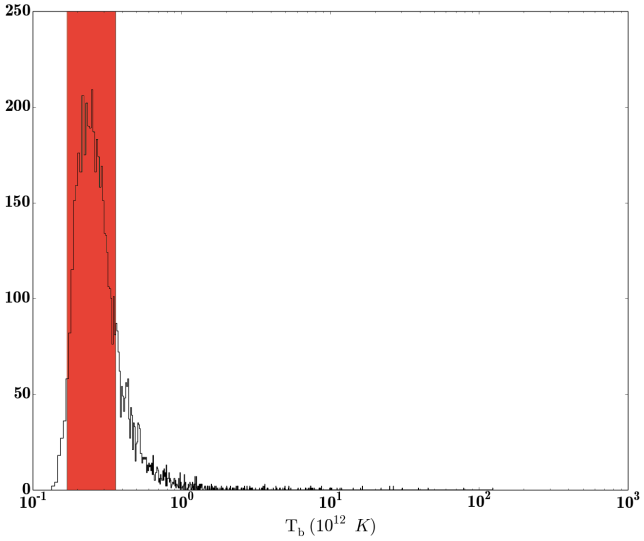

Figure 8. Histogram of the brightness temperature of J0809+5341 shown in black. The 68 per cent credible region (0.17 to 0.36) around the mode (0.25) is shaded light red.

δ is the *Doppler boosting factor*. Since $T_{b,int}$ is expected to lie below the inverse-Compton limit, our result implies that the source is indeed relativistically Doppler boosted i.e., $\delta > 1$ (see also An et al. (2016)).

6 A BAYESIAN CRITERION FOR THE RESOLUTION LIMIT

Martí-Vidal et al. (2012) derived the maximum theoretical over-resolution power of an interferometer, dependent on the signal-to-noise ratio of the visibility measurements, from theoretical considerations. The minimum resolvable source size θ_M is given by

$$\theta_M = \beta \left(\frac{l_c}{2(\mathrm{SNR})^2} \right)^{1/4} \times \mathrm{FWHM} \quad . \tag{28}$$

β depends on the shape and the intensity profile of the source model and takes values between 0.5 and 1. For source profiles with higher intensities at lower scales, the value of β is closer to 1. l_c is the log-likelihood value that corresponds to the critical probability of the null hypothesis, which is taken to be the point source model. It takes the values 3.84 and 8.81 for a 5 per cent and a 0.3 per cent probability cut-off of the null hypothesis respectively (Martí-Vidal et al. 2012). The FWHM is the *full-width at half maximum* of a circular PSF. For an interferometer with an elliptical PSF such as ours, the circular FWHM equivalent is given by $\sqrt{ab}$, where a and b are the major and minor axes of the PSF respectively (Lobanov 2005).

The SNR is calculated by computing the ratio between the weighted average of the visibilities and the noise $\sigma = \sigma_{\mathrm{vis}}/\sqrt{N}$. If σ varies between measurements (equation 16), then by *Parseval's theorem*[19], the rms noise σ_{rms} in the naturally-weighted residual image of the sky (after the source has been subtracted out) can be equated to σ as

$$\sigma^2 = \frac{\sigma^2_{\mathrm{vis}}}{N} = \sigma^2_{\mathrm{rms}} \quad , \tag{29}$$
$$\sigma = \sigma_{\mathrm{rms}} \quad .$$

In the Bayesian approach, we set the minimum resolvable size, θ_B, to the size of the source at which the evidence for CIRC against PT turns *positive* (Table 1). To compare θ_M with θ_B, we simulated a series of observations, each with a compact circular Gaussian source of a different size located at the pointing centre of the interferometer. For the station gains and SEFDs, we used the *maximum a posteriori* (MAP) estimates[20] obtained from analysing the actual data. We analysed three such sets of simulated data for three different SNR levels, where SNR is calculated according to equation (29), and computed the Bayes factor between CIRC and PT (Table 6).

Equation (28) gives an estimated resolution limit assuming perfect calibration. Martí-Vidal et al. (2012) observe how various factors such as the proportion of long baselines in an array configuration, source structure, and biased gains of the antennas providing the long baselines will further limit the resolving power of the interferometer. The Bayesian approach is sensitive to factors such as source shape, biased instrumental gains, and the associated uncertainties and provides a more realistic estimate of the minimum resolvable source size.

Table 6. A comparison of the resolution limits denoted by θ_M and θ_B, obtained using equation (28) with $\beta = 1$ and $l_c = 8.81$ and by using the Bayesian approach respectively, for the sparse VLBI array described in Table 2 for different SNR levels at 5 GHz. The circular FWHM equivalent of the naturally-weighted elliptical PSF is $\sqrt{5.7 \times 2.2} = 3.54$ mas. For higher SNR levels, θ_B is limited more by the gain amplitude calibration than by the theoretical capabilities of the array. The last column gives the maximum brightness temperature one could measure for a source of 1 Jy, derived from θ_B.

SNR	θ_M (mas)	θ_B (mas)	θ_B/FWHM	T_b/S_ν (10^{12} K/Jy)
150	0.42	0.45	0.13	0.763
2000	0.11	0.29	0.09	1.837
5900	0.07	0.17	0.05	5.347

7 CONCLUSION

Based on the Bayes factors obtained, we have very strong evidence for a resolved source with slightly elongated structure, in agreement with what is expected for a partially synchrotron self-absorbed compact jet in a flaring blazar. By simultaneously estimating source parameters along with the antenna gains and SEFDs, we have also acquired knowledge of the precision of our estimate of the source size and its correlation with antenna gains. Without SH visibilities, the Bayesian evidence indicates that the source must be considered unresolved.

$|g_{\mathrm{SH}}|$ varies by about 12 per cent, which is about an order of

[19] Parseval's theorem, in this context, ensures that the total power of the noise remains the same regardless of whether it is computed in the image domain or in the *uv*-domain.

[20] A MAP estimate is the Bayesian equivalent of the Maximum-Likelihood (ML) estimate and corresponds to a mode of the corresponding posterior distribution (Sivia & Skilling 2006). It may be seen as a regularised ML estimate.

magnitude worse than the precision in the gain amplitude calibration of all but one (NT) European station. The dependence of the shape parameters on $|g_{\rm SH}|$ results in the minor axis and the position angle of the Gaussian in model GAU being poorly constrained. This also illustrates the necessity of accounting for calibration errors for the stations providing the long baselines, lest we are led astray in our attempts to estimate the parameters relevant for the science goals.

Currently, we are limited only by the performance of the MultiNest sampler required to compute the evidence. The RIME can model any time or frequency variation in the source and in the instrumental effects. Within this framework, data from multiple spectral bands and epochs can be analysed together, incorporating time and frequency variation in the complex antenna gains. Future analyses will benefit from numerical samplers for evidence computation tailored for higher-dimensional parameter spaces.

The brightness temperature distribution we have derived for J0809+5341 accounts for the uncertainties in the source shape and instrumental gain calibration and indicates that the source is Doppler boosted and that the intrinsic brightness temperature must be less than what we obtain, consistent with the literature.

The analysis of synthetic observations of different SNR levels shows that, as the SNR improves (SNR$\gg$100), we are constrained less by the theoretical capabilities of the interferometer array and more by the effects of miscalibration of station gains and the discrepancy between θ_M and θ_B becomes more pronounced. The constraints derived by Martí-Vidal et al. (2012) assume perfect calibration. Our approach is more sensitive to uncertainties in source shape estimates and their correlation with instrumental gains and provides more realistic estimates that are 2 to 2.5 times the estimates obtained from equation (28). Knowledge of θ_B also enables us to estimate the maximum brightness temperature that one can measure for a given interferometer configuration. For sensitive future VLBI arrays such as the SKA-VLBI, it is possible to resolve source structure down to, or even less than, about 5 per cent of the size of the naturally weighted restoring beam with high precision, if the gain calibration of the stations providing the longest baselines is precise to within 1 per cent.

ACKNOWLEDGEMENTS

The EVN is a joint facility of independent European, African, Asian, and North American radio astronomy institutes. The data presented in this publication are derived from the EVN project code RSC02. We are grateful to An Tao and his group for sharing these data with us and to Ivan Martí-Vidal for his helpful comments. We also thank the referee for their comments. Part of the research leading to these results has received funding from the European Commission Seventh Framework Programme (FP/2007-2013) under grant agreement No. 283393 (RadioNet3).

IN and KH acknowledge funding from the MeerKAT High Performance Computing for Radio Astronomy Programme. JZ and IN acknowledge the South Africa National Research Foundation Square Kilometre Array Project for financial support. OS is supported by the South African Research Chairs Initiative of the Department of Science and Technology and National Research Foundation. Most of the computations leading to these results were carried out on Rhodes University's RATT (Radio Astronomy Techniques & Technologies) servers.

REFERENCES

Ables J. G., 1974, Astronomy & Astrophysics, 15, 383
An T., Cui Y.-Z., Paragi Z., Frey S., Gurvits L. I., Gabányi K. É., 2016, Publications of the Astronomical Society of Japan, 68, 77
Antonucci R., 1993, Annual Reviews of Astronomy & Astrophysics, 31, 473
Bertero M., de Mol C., 1996, Progress in Optics, XXXVI, 129
Born M., Wolf E., 1999, Principles of Optics, 7th edn. Cambridge University Press
Briggs D., 1995, PhD thesis, New Mexico Institute of Mining and Technology
Buchner J., et al., 2014, Astronomy & Astrophysics, 564, A125
Feroz F., Hobson M. P., 2008, Monthly Notices of the Royal Astronomical Society, 384, 449
Feroz F., Hobson M. P., Bridges M., 2009, Monthly Notices of the Royal Astronomical Society, 398, 1601
Feroz F., Hobson M. P., Cameron E., Pettitt A. N., 2013, astro-ph/1306.2144v2
Goodman J., 1968, Introduction to Fourier Optics, 1st edn. McGraw-Hill Book Company
Hamaker J. P., Bregman J. D., Sault R. J., 1996, Astronomy & Astrophysics, 117, 137
Harris J. L., 1964, Journal of the Optical Society of America, 54, 931
Hastings W. K., 1970, Biometrika, 57, 97
Högbom J. A., 1974a, Astronomy & Astrophysics, 15, 417
Högbom J. A., 1974b, Astronomy & Astrophysics Supplement Series, 15, 417
Honma M., Akiyama K., Uemura M., Ikeda S., 2014, arXiv e-prints, astro-ph/1407.2422,
Jaynes E. T., 2003, Probability Theory: The Logic of Science, 1st edn. Cambridge University Press
Jeffreys H., 1961, Theory of Probability, 3rd edn. Clarendon Press
Kass R. E., Raftery A. E., 1995, Journal of the American Statistical Association, 90, 773
Kawaguchi N., Jiang W., Shen Z.-Q., 2015, PASJ, 67, 112
Kellerman K., Owen F., 1988, Galactic and Extragalactic Radio Astronomy, p. 653
Kellerman K. I., Pauliny-Toth I. I. K., 1969, The Astrophysical Journal, 155, L71
Klöckner A., Pinto N., Lee Y., Catanzaro B., Ivanov P., Fasih A., 2012, Parallel Computing, 38 (3), 157
Lobanov A. P., 2005, arXiv e-prints, astro-ph/0503225,
Lochner M., Natarajan I., Zwart J., Smirnov O., Bassett B., Oozeer N., Kunz M., 2015, Monthly Notices of the Royal Astronomical Society, 450, 1308
Mackay D. J. C., 2003, Information Theory, Inference, and Learning Algorithms, 1st edn. Cambridge University Press
Marscher A. P., Scott J. S., 1980, Publications of the Astronomical Society of the Pacific, 92, 127
Martí-Vidal I., Pérez-Torres M. A., Lobanov A. P., 2012, Astronomy & Astrophysics, 541, A135
Martí-Vidal I., Vlemmings W. H. T., Muller S., Casey S., 2014, Astronomy & Astrophysics, 563, A136
Metropolis N., Rosenbluth A., Rosenbluth M., Teller A., Teller E., 1953, The Journal of Chemical Physics, 21, 1087
Middelberg E., Bach U., 2008, Reports on Progress in Physics, 71, 6901
Noordam J. E., Smirnov O. M., 2010, Astronomy & Astrophysics, 524, A61
Paragi Z., et al., 2015, Advancing Astrophysics with the Square Kilometre Array (AASKA14), p. 143
Pâris I., et al., 2014, Astronomy & Astrophysics, 563, A54
Parker R. L., 1977, Annual Review of Earth and Planetary Sciences, 5, 35
Perkins S., Marais P., Zwart J., Natarajan I., Tasse C., Smirnov O., 2015, Astronomy and Computing, 12, 73
Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 2007, Numerical Recipes: The Art of Scientific Computing, 3rd edn. Cambridge University Press
Rayleigh 1879, Philosophical Magazine and Journal of Science, 8, 49

Readhead A. C. S., Wilkinson P. N., 1978, The Astrophysical Journal, 223, 25
Rees M. J., 1966, Nature, 211, 468
Rybicki G., Lightman A., 1985, Radiative Processes in Astrophysics. Wiley-VCH
Ryle M., 1955, Vistas in Astronomy, 1, 532
Schwab F. R., 1984, The Astronomical Journal, 89, 1076
Shepherd M. C., 1997, Astronomical Data Analysis Software and Systems VI, ASP Conference Series, 125
Sivia D. S., Skilling J., 2006, Data Analysis: A Bayesian Tutorial, 2nd edn. Oxford University Press
Skilling J., 2004, in R. F., R. P., V. T. U., eds, 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering. pp 395–405
Smirnov O. M., 2011, Astronomy & Astrophysics, 527, A106
Tarantola A., 2005, Inverse Problem Theory and Methods for Model Parameter Estimation, 1st edn. SIAM
Taylor G. B., Carilli C. L., Perley R. A., eds, 1999, Synthesis Imaging in Radio Astronomy II, 1st edn. Astronomical Society of the Pacific Conference Series Vol. 180, Astronomical Society of the Pacific Conference Series
Thompson A. R., Moran J. M., George W. Swenson J., 2001, Interferometry and Synthesis in Radio Astronomy, 2nd edn. Wiley-Interscience
Trotta R., 2008, Contemporary Physics, 49, 71
Urry C. M., Padovani P., 1995, Publications of the Astronomical Society of the Pacific, 107, 803
Wilson T., Rohlfs K., Hüttemeister S., 2009, Tools of Radio Astronomy, 5th edn. Springer
di Francia G. T., 1969, Journal of the Optical Society of America, 59, 799

This paper has been typeset from a TEX/LATEX file prepared by the author.